

Understanding how
platforms with video-
sharing capabilities
protect users from
harmful content online

August 2021

contents

02

Executive
summary

05

Background,
scope and
methodology

09

Current state
assessment

16

Measures and
costs

35

User experience





This report has been prepared by Ernst & Young LLP, a limited liability partnership registered in England and Wales with registered number OC300001, in accordance with an engagement agreement for professional services with DCMS. Ernst & Young LLP's obligations to DCMS are governed by that engagement agreement. This disclaimer applies to all other parties (including DCMS' affiliates and advisors).

This report has been prepared for general informational purposes only and is not intended to be relied upon as accounting, tax, or other professional advice. Refer to your advisors for specific advice.

Ernst & Young LLP accepts no responsibility to update this report in light of subsequent events or for any other reason.

This report does not constitute a recommendation or endorsement by Ernst & Young LLP to invest in, sell, or otherwise use any of the markets or companies referred to in it.

To the fullest extent permitted by law, Ernst & Young LLP and its members, employees and agents do not accept or assume any responsibility or liability in respect of this report, or decisions based on it, to any reader of the report. Should such readers choose to rely on this report, then they do so at their own risk.

Ernst & Young LLP reserves all rights in the Report.

1

Executive summary

Scope and context

In this report, we assess platforms with video-sharing capabilities, a type of online service that allow users to upload and share videos. Some of these platforms are primarily focused on sharing user-generated videos; others are social media sites that allow users to upload videos as part of a wider service.

We have analysed:

- ▶ Growth, competition and innovation in the sector.
- ▶ The measures these platforms take to protect their users online.
- ▶ The costs of these measures.
- ▶ The user experience on these platforms.

Our findings are based on our analysis of data on UK consumption of online video content, interviews with seven platforms with video-sharing capabilities, a survey of 12 platforms with video-sharing capabilities, and consumer research. This report is intended to provide a snapshot of the sector at a point in time and to give a sense of the measures that platforms take to protect their users and the associated costs.



Current state assessment of platforms in the UK with video sharing capabilities

Our analysis has found that users are consuming more content now than ever before. In 2020, UK audiences watched over 213 billion videos. This represents an increase of 40 billion in the number of videos being watched online since 2017, or average growth of 5% per year over the same period. The average time users over 18 spend watching videos across all platforms has grown to 32 minutes per week per user, an increase of 8% per year since 2017. This average includes the time spent on all platforms within our analysis, including YouTube. The average time spent by over 18s per week per user watching videos on YouTube is five and a half hours – this is significantly higher than the overall average across all platforms.

The sector is highly concentrated, with consumers using a limited number of platforms to view videos online. Only a small number of platforms have entered the sector and achieved scale in recent years, for example TikTok, which launched in the UK in 2017. Our analysis shows that new entrants since December 2017 are adult (pornographic) sites. Of the new adult sites identified, none have achieved scale.

Platforms that enter the market are faced with several barriers to compete effectively with established platforms. For example, established platforms can expand their product offerings in response to the launch of new platforms and can acquire smaller platforms to enhance the platform's ecosystem of apps. In general, successful new entrants have provided an innovative product offering, differentiating themselves from established platforms.

213 billion videos viewed online in the UK in 2020

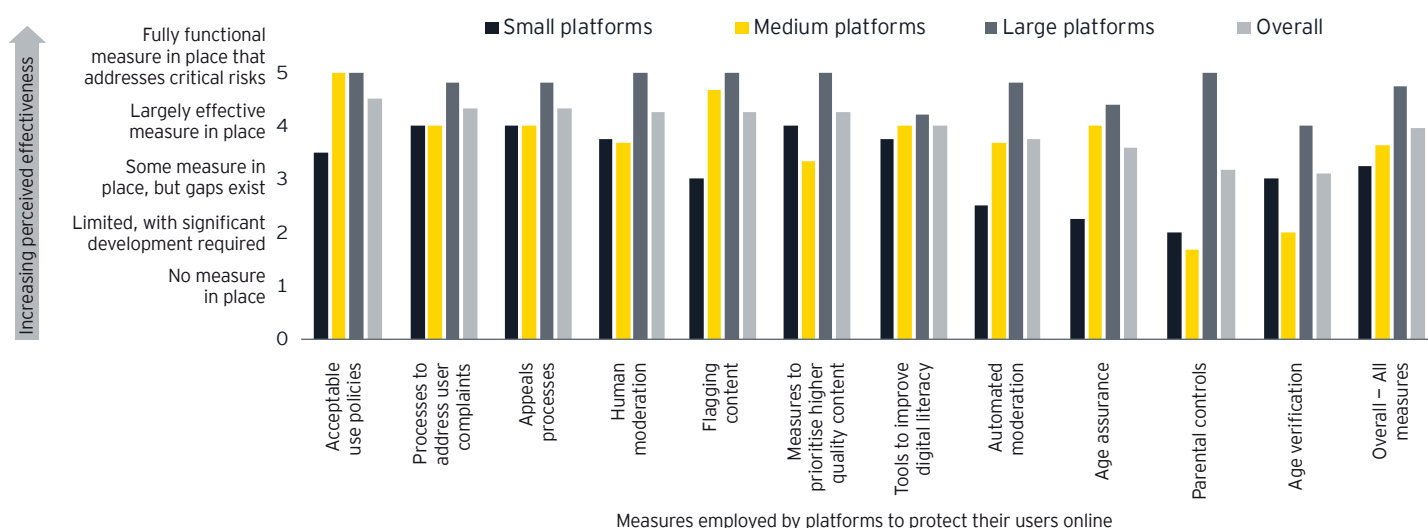
5% annual growth in the total number of videos viewed online in the UK since 2017

Average time viewing videos online of **32 minutes** per week across all platforms

Analysis of the measures that platforms take to protect users online and the associated costs

Our analysis shows that different platforms take different approaches to protecting users online, driven by factors including the size of the platform, the resources it has to invest in online safety, its assessment of the risks specific to the platform, public pressure and, in some cases, competitive considerations. Our survey suggested, in general, large platforms¹ have the most confidence in the measures they have in place to protect users online. With some exceptions, small platforms are more likely to state that the measures they have in place require significant development to effectively protect users online.

Platforms' perceived effectiveness of measures to protect users online by platform size



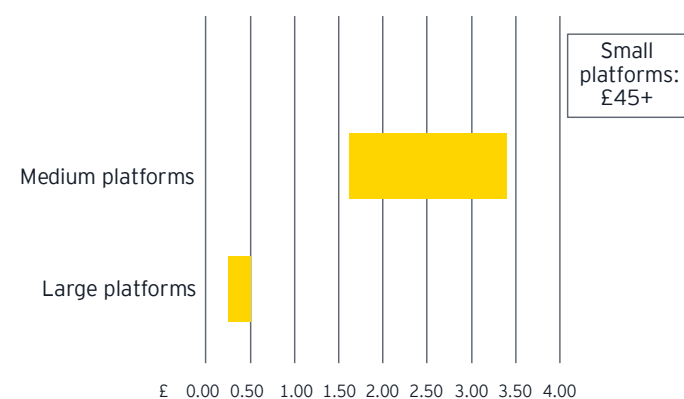
Annual expenditure on measures to protect users online varies significantly by platform size. Total annual expenditure ranged from hundreds of pounds for the very smallest platforms to over £1.5bn for the largest platforms included in our research.

Large platforms benefit from significant economies of scale when implementing online safety measures. The cost per user ranged from £0.25 to £0.50 for large platforms included in our research, with evidence suggesting costs for some of the largest platforms may be materially lower. In comparison, small platforms included in our research (some of which may be loss-making start-ups) spent over £45 per user on online safety measures.

We compared the perceived effectiveness of measures against the annual amount invested in online safety for each platform included in our survey. We found the perceived effectiveness of measures increases with increased investment. Our analysis also suggests that significant incremental investment is required to make marginal improvements to measures.

Most platforms included in our research suggested the implementation of the EU's Audiovisual Media Service Directive, which requires platforms with video-sharing capabilities to take certain measures to protect their users online, would not lead to them spending more on online safety measures than initially anticipated as part of their internal strategies.²

Cost per user by platform size



¹ We refer to 'small' platforms as those having less than 100k unique global users, 'medium' platforms as those having between 100k and 10m unique global users, and 'large' platforms as having more than 10m unique global users.

² Further information on the EU's Audiovisual Media Service Directive can be found on page 5.

The user experience of platforms with video-sharing capabilities

Our research suggests that the way users engage with platforms with video-sharing capabilities means there is a risk of accidental exposure to harmful content online: recommendation algorithms play an important role helping people find videos to watch, and one of the most common approaches to deciding whether to watch a video is to make a judgement once the video has started playing. However, the online safety measures that we have reviewed as part of this report, set out in detail in section 4, have limited focus on recommendation algorithms, despite this being a common way in which video content is accessed.

Additionally, the way that children use platforms with video-sharing capabilities suggests that children may be particularly at risk of being exposed to harmful content online: most young children watch videos without supervision and are attracted to sites designed for older children. Platforms therefore need to consider the potential for their content to be accessed by children regardless of whether their content is intended for children or not.

One in four adults say they browse videos recommended by the platform itself

45% of adults make a judgement on the appropriateness of the video after they have started watching it

The platforms included in our research reported having largely effective measures in place to protect users from harmful content online. However, our consumer research showed the scale of harmful content accessed online among all age groups, with exposure to harmful content on platforms with video-sharing capabilities particularly pronounced among under 18s. Our analysis suggests that significant incremental investment would be required to materially improve the effectiveness of measures employed by platforms.

A third of adults and half of 13-17 year olds have seen content on platforms with video-sharing capabilities in the past three months that they think may have been illegal (but they are unsure)

78% of 6-12 year olds have experienced some type of harmful content on platforms with video-sharing capabilities

People have differing attitudes to harmful content online and the appropriate response to such content. Our research highlights confusion among some users as to the role of platforms, the rights of users and the extent to which platforms and government should intervene. The measures employed by platforms need to balance the expectations and priorities of different users.

19% of UK adults agree with the statement: "It's not right that video-sharing platforms are increasingly policing/censoring content that is shared on their sites", while 51% disagree and 30% neither agree nor disagree

59% of UK adults agree with the statement: "Video-sharing platforms are not doing enough to keep illegal content off their platforms"

Many people do not engage with information on platforms' websites, and those that do engage struggle to understand the information available. A large proportion of respondents to our survey were resistant to safety measures that introduce friction in the user experience, suggesting that interventions that are integrated in a way that minimises friction would be most effective in protecting users online.

39% of UK adults agree with the statement: "I only look at the videos on video-sharing platforms, so wouldn't pay attention to any additional links or information on the site designed to stop exposure to online harm"

42% of UK adults agree with the statement: "I prefer video platforms where there is no barrier between me and the video if I click on it to watch"

2

Online platforms are increasingly popular but historically have not been regulated

While content on other platforms is regulated, there has been limited regulation of online platforms

Television programmes are required to comply with Ofcom's Broadcasting Code. The Broadcasting Code sets out guidance for broadcasters on safeguards to protect children from harmful content. Among other issues, the Code also provides extensive guidelines on how broadcasters should handle programmes that may cause harm and offence, and programmes that include content related to crime, disorder, hatred and abuse. Television advertising is also regulated.

Ofcom also regulates video-on-demand platforms, such as those providing catch up TV and a library of archive content, though to a lesser extent than programmes broadcast on television. Video-on-demand platforms must protect viewers, particularly children, from harmful and illegal content, and there are also rules governing advertising.

Historically, other types of online platforms have not been regulated. However, there are growing concerns about harmful and illegal content online, and particularly about the exposure of children to harmful online content.

New regulation will protect users from harmful and illegal content appearing on online platforms

The EU's Audiovisual Media Services Directive (AVMSD) is the regulatory framework that governs Member States' national legislation on audiovisual media. Recent updates to the AVMSD have extended the scope of regulation to video-sharing platforms (VSPs) for the first time. The rules that govern what constitutes a VSP are complex, though they are broadly a type of online service that allow users to upload and share videos. The AVMSD, as transposed into domestic UK law, now requires VSP providers to take appropriate measures to protect under-18s from



potentially harmful material and to protect the general public from incitement to hatred or violence and other specific material the inclusion of which would be a criminal offence. Services also need to ensure certain standards around advertising are met. The regulatory framework was transposed into UK law in September 2020 and came into force on 1 November 2020. Ofcom is the UK's national regulatory authority for the VSP regime.

Separately, the UK government has set out its intention to introduce the Online Safety regime under which a new duty of care will be placed on online services towards their users. The Online Safety regime shares broadly similar objectives to the VSP regime, though the scope of the Online Safety regime is broader, both in terms of the companies in scope, but also in terms of the regulatory framework. The government intends for requirements on UK-established VSPs to be superseded by the Online Safety regime, once the latter comes into force. The Online Safety Bill, which will introduce the Online Safety framework, has now been published in draft.³

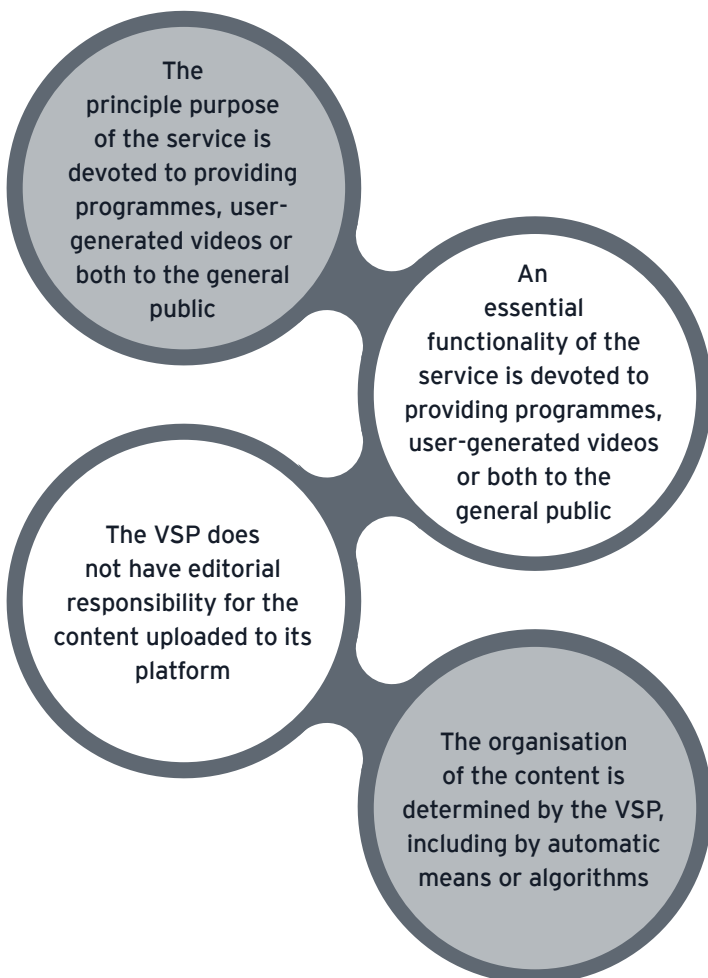
³ <https://www.gov.uk/government/publications/draft-online-safety-bill>

This report includes information about online platforms that allow users to upload and share videos

What is a video-sharing platform?

Video-sharing platforms (VSPs) are a type of online service that allow users to upload and share videos. Some VSPs are primarily focused on sharing user-generated videos; others are social media sites that allow users to upload videos as part of a wider service.

Definition of a VSP under the AVMSD



Ofcom has published guidance to help VSPs assess whether their service is in scope for regulation

The AVMSD includes a detailed definition of VSPs subject to the new regulation, illustrated in the image to the left.

The AVMSD also sets out criteria for determining the jurisdiction of VSPs. Like in broadcasting, VSPs will fall under the jurisdiction of the national regulator in the country where the VSP is established. As such, only platforms deemed to be VSPs and with a UK establishment as decided by the applicable rules, will be regulated by Ofcom.

Only a small number of VSPs are likely to fall within Ofcom's scope of regulation. Ofcom has published guidance to help VSPs understand whether they fall within the scope of the AVMSD and will continue working closely with VSPs on this matter. VSPs are required to notify Ofcom if they believe they will be within the UK's jurisdiction. Notification for the VSP regime concluded on 6 May 2021.

These publications can all be found on Ofcom's website.

We have not assessed whether the platforms included in our research would be in scope for regulation

The platforms that we have included in our research all provide some capability for users to upload and share videos, without the platform exercising editorial control over the content uploaded.

We have not formed an opinion on whether any of the platforms included in our research would be in scope for any UK legislation or regulation that has or might come into effect in this area.

DCMS has commissioned EY to research the measures that platforms with video-sharing capabilities take to protect users online

While other countries have introduced legislation to address specific types of harm, the UK's Online Safety regime is the first attempt globally to address a spectrum of online harms through a single regulatory approach.

DCMS's objectives for the Online Safety regime include ensuring that the requirements are feasible and undue burdens are not placed on businesses, as well as ensuring people are kept safe online.

In this context, EY has been commissioned by DCMS to evaluate the measures that platforms with video-sharing capabilities currently take to protect their users, and the associated costs to platforms of these measures, ahead of the implementation of the VSPs regulatory regime and the later Online Safety regime.

Our research will ensure that costs to industry can be considered in policymaking.

The scope of our work is as follows:

- ▶ Carry out a current state assessment of platforms with video-sharing capabilities, analysing growth, competition and innovation in the sector, set out in section 3 of this report.
- ▶ Assess the measures that platforms with video-sharing capabilities have in place to protect users online, the associated costs of these measures, and how these costs vary by different types of platform, set out in section 4.
- ▶ Research user experience of platforms with video-sharing capabilities, including both children and adults, set out in section 5.




Our approach provides an initial view of how platforms with video-sharing capabilities protect their users and the costs they face

We have employed several research methods to gather information on how platforms with video-sharing capabilities protect their users online, the risks they face, and how sophisticated their measures are. To get a better understanding of what different platforms do and how the sector operates, we have:

- ▶ Analysed data on user numbers and viewing trends to better understand the evolution of platforms with video-sharing capabilities.
- ▶ Conducted a series of interviews with seven different online platforms and one provider of age assurance and age verification technologies to understand the measures platforms enforce, what types of cost they incur, and what barriers they face to protecting their users.
- ▶ Designed and administered a survey of 12 online platforms with video-sharing capabilities to better understand how sophisticated their measures are and the costs they incur to enforce those measures. The survey was conducted online and included platforms of varying sizes to gain a sense of variation across the sector.

- ▶ Undertaken desktop research to provide additional insight into user behaviour, online harms, and innovation in the sector.
- ▶ Designed and administered three separate consumer surveys focused on children, teenagers and adults to understand how they use platforms with video-sharing capabilities and their awareness of online harms. The sample sizes for the groups included in our survey are as follows: adults aged 18+ (1,000 respondents), children aged 13-17 (500 respondents), children aged 6-12 (500 respondents). These surveys were conducted online and targeted a representative sample of social media users.

Our analysis is intended to provide a snapshot of the sector at a point in time. It is not intended to provide a comprehensive review of the sector or detailed costing analysis, which would require more detailed data from platforms.



Presenting platforms of differing scales

For the purposes of explaining our findings, we refer to 'small' platforms as those having less than 100k unique global users, 'medium' platforms as those having between 100k and 10m unique global users, and 'large' platforms as having more than 10m unique global users.

3

We have analysed trends among platforms with video-sharing capabilities



We have carried out a current state assessment of trends in growth, competition and innovation among platforms with video-sharing capabilities. Our analysis covers the four-year period between January 2017 and December 2020. This period reflects the fast paced growth of the sector.

Our analysis focuses solely on the consumption of videos in this sector, and does not consider other usage trends for social media platforms over the time period. The data we have analysed relates solely to viewing videos rather than use of the platforms for any other purpose.

The basis of our current state assessment

The findings of our assessment are underpinned by quantitative and qualitative research, which includes analysis of key sector metrics, engagement with stakeholders in the sector (both platforms and their users), and a review of regulatory literature.

Some key sources of information include:

- ▶ **Comscore data:** We have used Comscore's Video Metrix Multi-Platform (VMX MP) dataset, which gives a view of digital video consumption across desktop, smartphones, tablets and over-the-top (OTT) media

devices, including in-app viewing. This data covers videos viewed by UK consumers only.

- ▶ **Literature review:** We have reviewed regulatory reports, including reports published by Ofcom and the Competition and Markets Authority.
- ▶ **Primary research:** We have conducted interviews with seven platforms with video-sharing capabilities and conducted research into the user experience of these platforms.

Growth



Since 2017, the total number of videos viewed on platforms with video-sharing capabilities in the UK has increased by an average of **5%** per year, with the average time spent by users increasing by **8%** per year to **32 minutes** per week

UK consumers watched 213 billion videos in 2020. The majority of these views were on YouTube (184 billion), which experienced average annual growth in the number of videos viewed of 6% per year over the period 2017 to 2020.

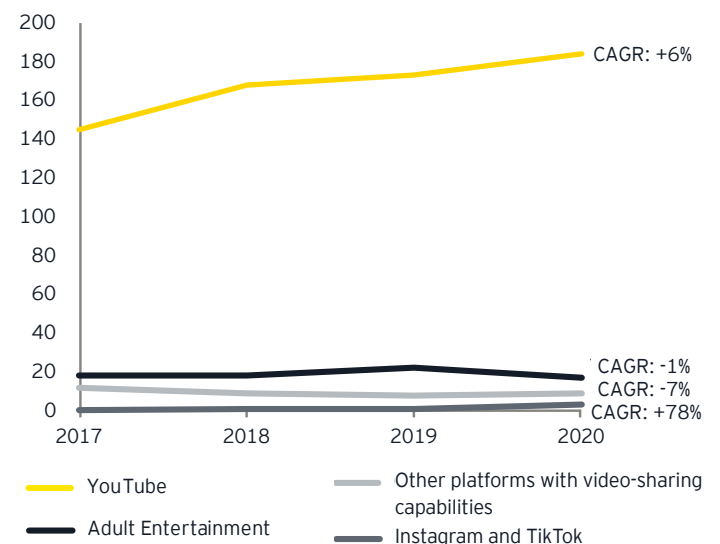
A small number of platforms, including Instagram and TikTok, have also experienced growth in terms of the number of users watching videos on the platform, the time they spend watching videos, and the number of videos they view. The number of videos watched in total on these two platforms has grown from c.300 million videos per year in 2017 to over 2.5bn by the end of 2020.

In contrast, adult (pornographic) entertainment sites have experienced a decline in video views, with the number of videos viewed decreasing by an average of 1.4% annually between 2017 and 2020, despite the number of adult sites in the sector increasing over the same time period.

In aggregate, other platforms included in our research (other platforms with video-sharing capabilities besides YouTube, adult entertainment sites, Instagram and TikTok) also experienced a decline. The number of videos viewed on these other platforms has decreased by an average of 7% a year since 2017.

Our analysis has focused only on video viewing, and not on other uses of these platforms.

Total annual video views in the UK (billion)



Source: Comscore VMX MP, EY custom-defined list of entities, total audience reach (over 18s), Jan 2017 – Dec 2020, United Kingdom.

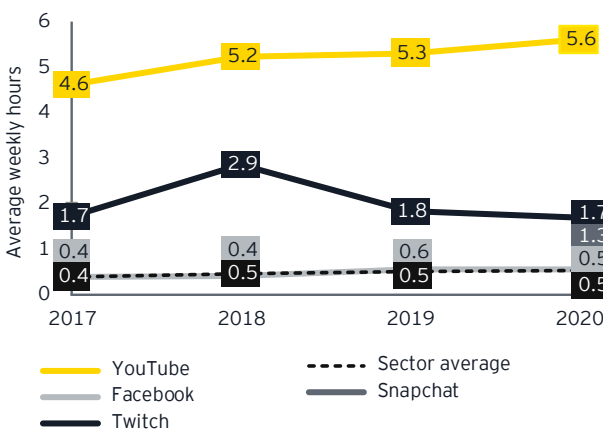
The average time users spend watching videos across all platforms has grown to 32 minutes per week, an increase of 8% per year or 8 minutes in total, since 2017.

The average time spent watching videos on YouTube is significantly higher, at over five and a half hours per week. Users are also spending more time watching videos on Facebook, increasing from 23 to 33 minutes per week since 2017.

The average time spent per user on YouTube is much higher than other platforms partially due to the wide variety of short-form and long-form content available. In 2020, the average time spent watching a single video on YouTube was 4 mins 12 seconds, compared to 1 min 48 seconds on Facebook.

In contrast, average time spent by users watching videos on Twitch has declined by 0.4% per year since 2017, although the number of users watching videos and the number of videos viewed have increased by 21% and 27% respectively. This suggests that the user base is growing, but that the time spent by new users on the platform has marginally declined.

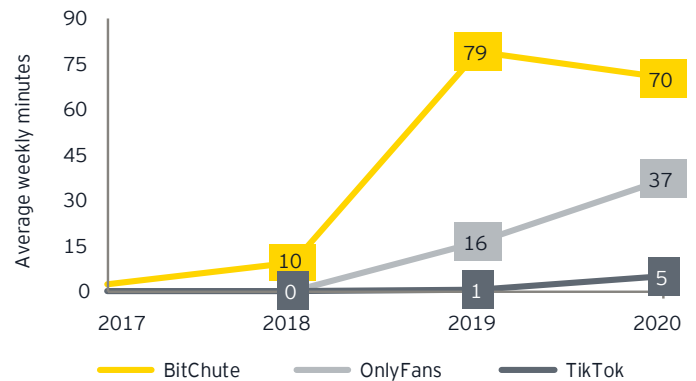
Average hours spent watching videos per week per user in the UK



Source: Comscore VMX MP, EY custom-defined list of entities, total audience minutes per viewer (over 18s), Jan 2017 – Dec 2020, United Kingdom. Snapchat figure (Sept 2019) taken from the Ofcom Online Nation (2020) Report.

Users are also spending an increasing amount of time watching videos on some newer-platforms. There has been a rise in the average time spent watching videos on BitChute (a video-hosting service), OnlyFans (a content subscription service) and TikTok (a video-sharing social networking service). Due to the short-form nature of TikTok videos, the average minutes spent viewing has not increased substantially, but the number of videos viewed on the platform has grown from an average of 1 video viewed per week per user in 2017, to 80 per week per user in 2020.

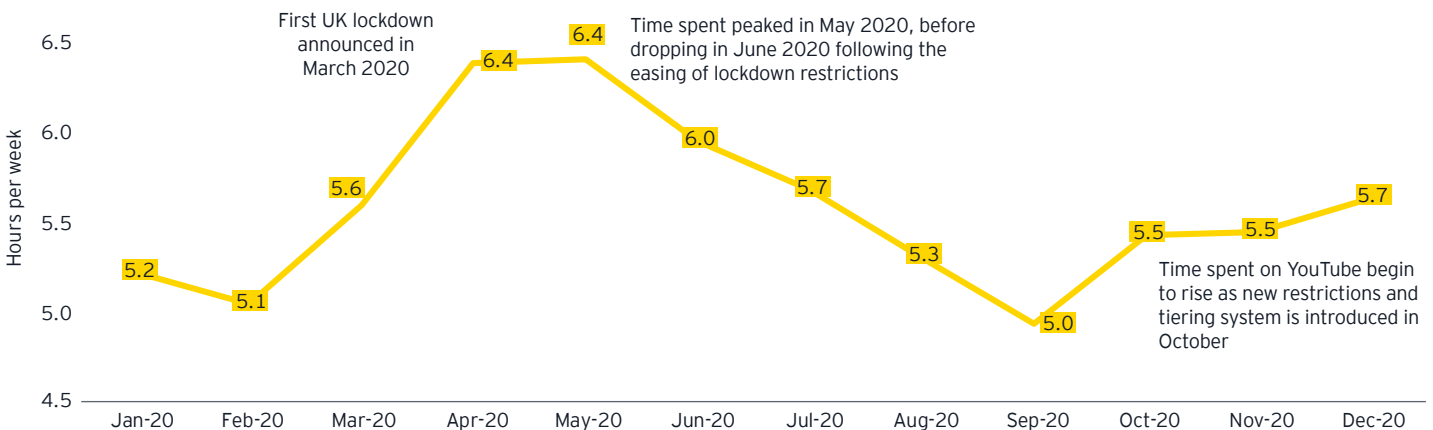
Average minutes spent watching videos per week per user in the UK



Source: Comscore VMX MP, EY custom-defined list of entities, total audience minutes per viewer (over 18s), Jan 2017 – Dec 2020, United Kingdom.

Throughout the COVID-19 pandemic, the amount of time people spend watching videos online has fluctuated with lockdown restrictions. After the first lockdown restrictions were announced in March 2020, average minutes spent watching videos online increased through April and May 2020 by 14%. When lockdown restrictions were eased in June 2020, the average time spent on YouTube decreased back to pre-lockdown levels, but it increased again towards the end of the year as restrictions were re-introduced.

Average hours spent watching YouTube videos per week per user in the UK throughout 2020



Source: Comscore VMX MP, total audience minutes per viewer (over 18s), Jan 2020 – Dec 2020, United Kingdom.

Competition

In general, consumers use a limited number of platforms to view videos online, and historically only a small number of platforms have achieved scale.

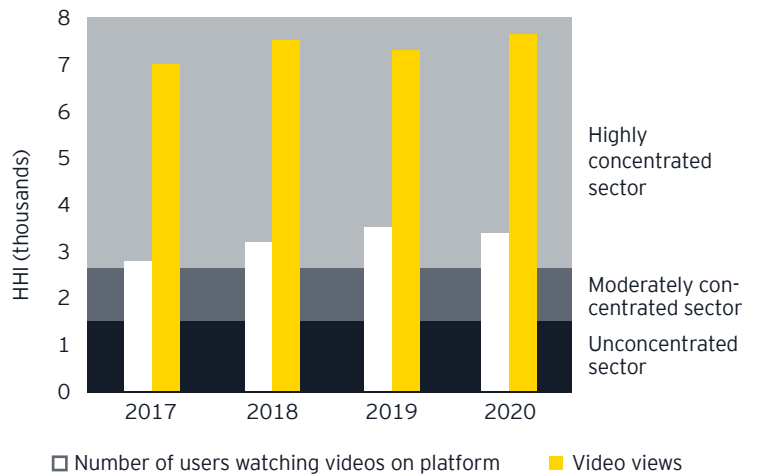


The Herfindahl-Hirschman Index (HHI) is used to assess the level of concentration in a sector.

HHI	What does this mean?
Less than 1,500	Unconcentrated sector
Of 1,500 to 2,500	Moderately concentrated sector
2,500 or greater	Highly concentrated sector

Based on our analysis of the number of users watching videos on platforms and the number of video views, the HHI results demonstrate that despite the entrance of some new platforms, the sector continues to be highly concentrated. When undertaking this analysis, we have focused only on video viewing, and not on other uses of social media platforms. The HHI results would differ if other uses of social media were included in the analysis.

HHI -UK sector concentration

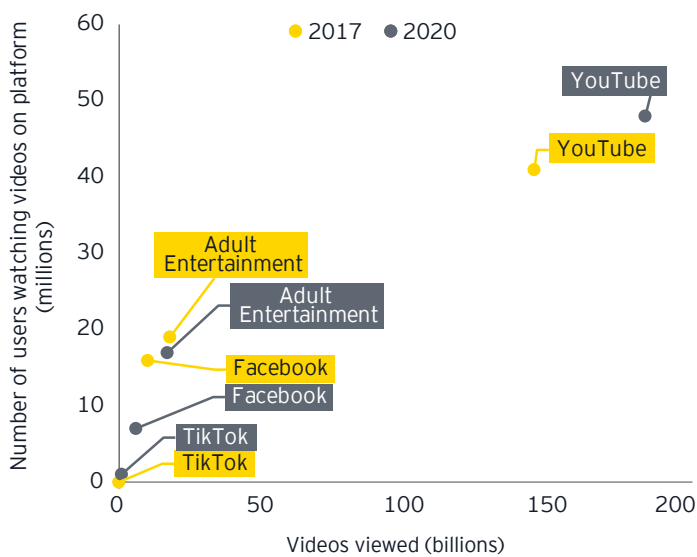


Source: Based on EY analysis of ComScore VMX MP data.

Our analysis suggests that people only use a small number of platforms to view videos. YouTube has experienced significant growth in both the number of people who use the platform to watch videos and the number of videos viewed on its platform between 2017 and 2020. In contrast, while Facebook has experienced growth in the number of videos viewed on the platform, the number of people using the platform to watch videos has remained stable.

78% of respondents to our consumer research said that YouTube and Facebook were the main platforms they use to watch videos online

Trends for selected platforms in the UK



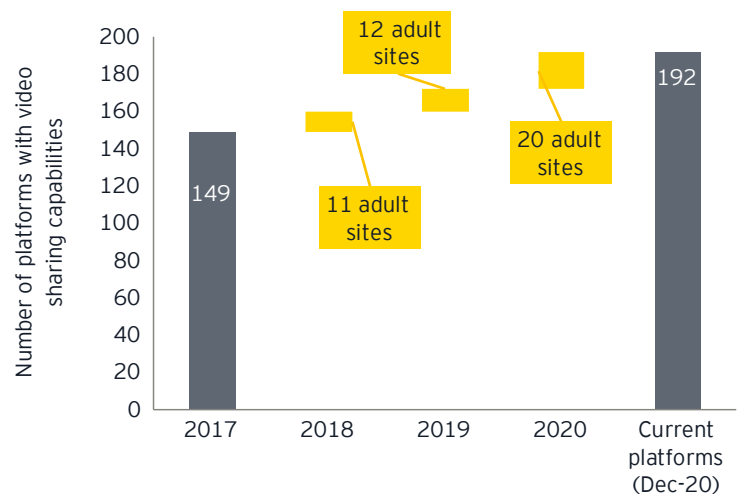
Source: Comscore VMX MP, total audience reach and total audience minutes per viewer (over 18s), Jan 2017-Dec 2020, United Kingdom.

Only a limited number of platforms have entered the sector and achieved relative scale in recent years, for example TikTok, which launched in the UK in 2017. Since December 2017, our analysis shows that new entrants to the sector have been adult sites, with 43 new adult sites launching. Of the new adult sites identified, none have achieved scale. It is not clear from our analysis why most new sites with video-sharing capabilities launched are adult sites.

Platforms that enter the market are faced with several barriers to be able to compete effectively with established platforms. Examples of the barriers faced include:

- ▶ Established platforms have the ability to expand their product offerings in response to the launch of new platforms, as a way to retain users.
- ▶ Established platforms can acquire smaller platforms to enhance the platform's 'ecosystem' of apps.
- ▶ According to our consumer research, 38% of people said the main reason they use a platform to watch videos is because it was the first platform they used, suggesting consumer inertia exists.

Number of platforms with video-sharing capabilities launched and achieving a minimum of 10 views per month in the UK



Source: Comscore VMX MP, total audience reach (over 18s), Jan 2017-Dec 2020, United Kingdom. This analysis is based on the point in time from which platform use became reportable on the Comscore platform.

Innovation



Differentiated experiences and features

In general, new entrants have only achieved scale where they have been able to offer users highly differentiated experiences and features. Examples of this include TikTok’s focus on short-form videos and the success of Snapchat’s ‘Stories’ feature, which launched in 2016. Similar ‘Stories’ features are now widely available on other social media sites.

Livestreaming and live chat

In recent years, livestreaming and live chat have also grown in popularity with consumers. Twitch has experienced significant growth in total number of video views as it enables gaming communities to stream and socialise, and exposes users to new games.

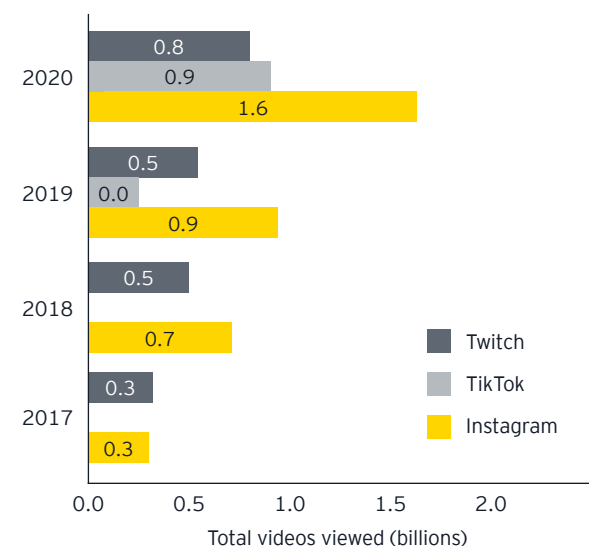
Expanded original content offering

Some platforms have expanded their offering of original content to attract and retain users. YouTube Originals offers original series and movies, and Snap Originals offers shows and interactive experiences. OnlyFans offers paying subscribers exclusive access to private content that isn’t available on other sites. The role of social media influencers is also a key consideration for platforms, as influencers can attract consumers to platforms with their content.

Platforms curate content to arrange how videos become visible to platform viewers

The curation of content in a way that is appealing to users can maximise viewer engagement with the platform. Simple search functions can help to attract and retain users. Additionally, platforms are increasingly focusing on tailoring recommendations to ensure users see content that appeals to them, and also to reduce ‘search fatigue’ among users who want to spend less time searching for content.

Total videos viewed on selected platforms in the UK



Source: Comscore VMX MP, total audience videos viewed, Jan 2017 – Dec 2020, United Kingdom.

Findings

Current state assessment

1

Consumers in the UK are watching more videos than ever before through online platforms, with over 213 billion videos watched in 2020. This represents average annual growth of 5% since 2017. YouTube remains the most popular platform for watching videos online.

2

The average time consumers spend watching videos is growing. YouTube has the highest average time spent watching videos per week (5.6hrs), with other platforms on average seeing an increase of 8% per year in time spent watching videos (which is reflective of an increase of 9 minutes) since 2017.

3

The sector is highly concentrated, with consumers using a limited number of platforms to view videos online. Only a small number of new platforms have entered the sector and achieved scale in recent years. Our analysis shows that new entrants since 2017 have been adult sites.

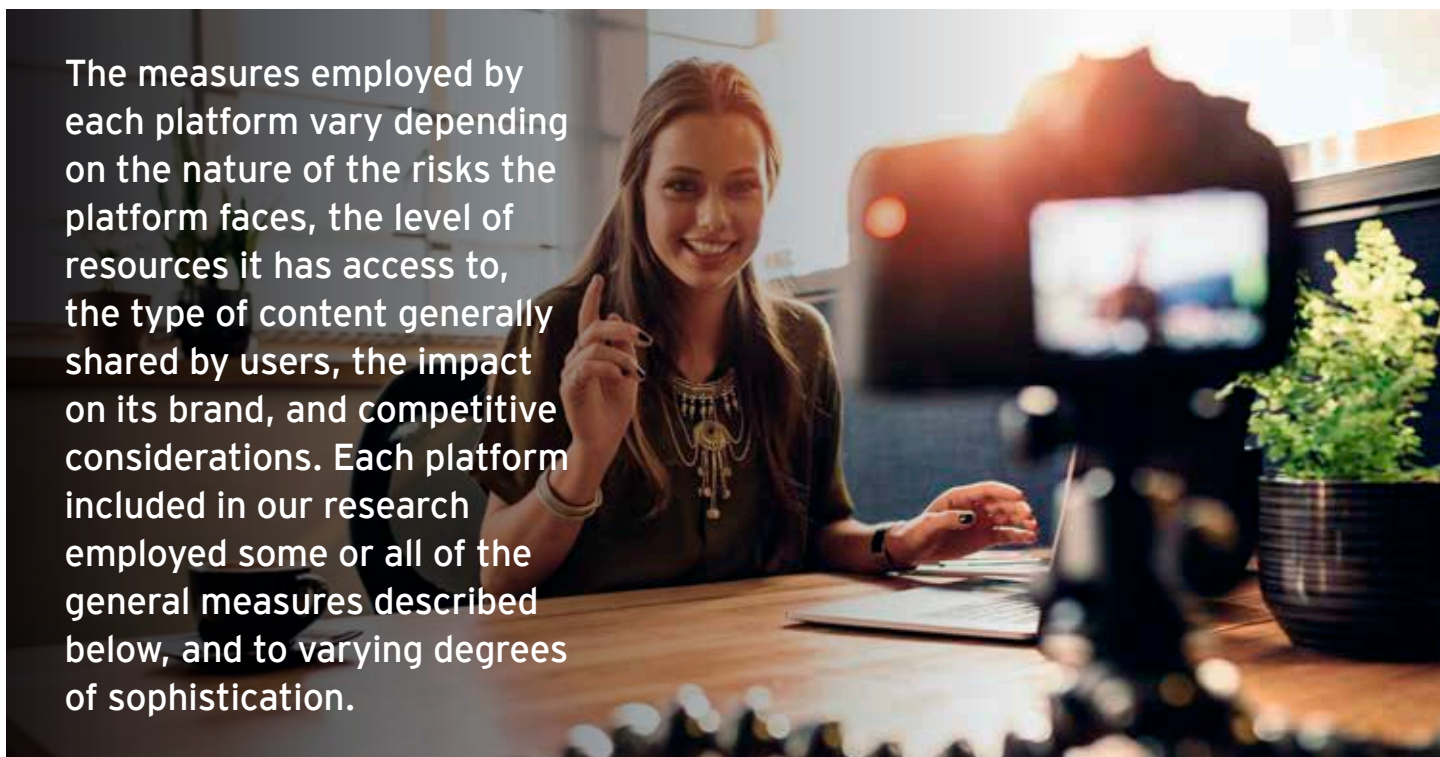
4

New platforms launching are faced with several barriers to compete effectively with established platforms. In general, successful new entrants have provided an innovative offering, differentiating themselves from established platforms. Established platforms have expanded their product offerings to remain competitive and relevant to their users.

4

Different measures are employed by different platforms to protect their users from online harms

The measures employed by each platform vary depending on the nature of the risks the platform faces, the level of resources it has access to, the type of content generally shared by users, the impact on its brand, and competitive considerations. Each platform included in our research employed some or all of the general measures described below, and to varying degrees of sophistication.



Policy

General terms and conditions: Set out principles that users must adhere to

Acceptable use policies: Set out types of behaviour that is considered unacceptable on platforms.

Community guidelines: Set out specific rules for sub-sections of platforms that must be adhered to.

Prevention

Age assurance and age verification: Assurance involves estimating a user's age using technology, and age verification involves corroborating a user's age against photo ID or credit card details.

Pre-upload review: Reviewing material before it is made available for public consumption on a platform.

Parental controls: Blocking access to adult websites or limiting the functionality of apps to make them suitable for children.

Prompts: Reminding users of platform policies and making users agree to these policies before uploading new content.

Detection

Moderation: Automated moderation involves implementing algorithms and software to review and remove potentially harmful content. Human moderation involves physical review of material by a person.

Flagging: Functionality for all users to flag potentially harmful content for review. Trusted flagging is when content is prioritised for review as it is flagged by a trusted user.

Review

Appeals process: A system in place to allow users to challenge actions taken by platforms including content removal.

Transparency reporting: Tracking specific transparency metrics to provide insight into the level of potentially harmful or inappropriate content being removed from a platform, and what impact that has on its users

Child protection

There are several industry wide resources used to prevent the spread of Child Sexual Abuse Material (CSAM) and Child Sexual Abuse Imagery (CSAI):

Use of hashed (tagged) images: Microsoft and YouTube have created databases of 'hashed' (tagged) images of inappropriate material. They offer these databases to platforms to check material against so it can be quickly removed.

Use of scanning software: Google has created software which scans new material for potential CSAM and CSAI to prioritise it for moderation.

Working with non-profit organisations: Most platforms we spoke to also worked with the National Centre for Missing or Exploited Children (NCMEC) or the Internet Watch Foundation (IWF) to report instances of possible child exploitation.

Overall, large platforms have the most confidence in the measures they have in place to protect users online

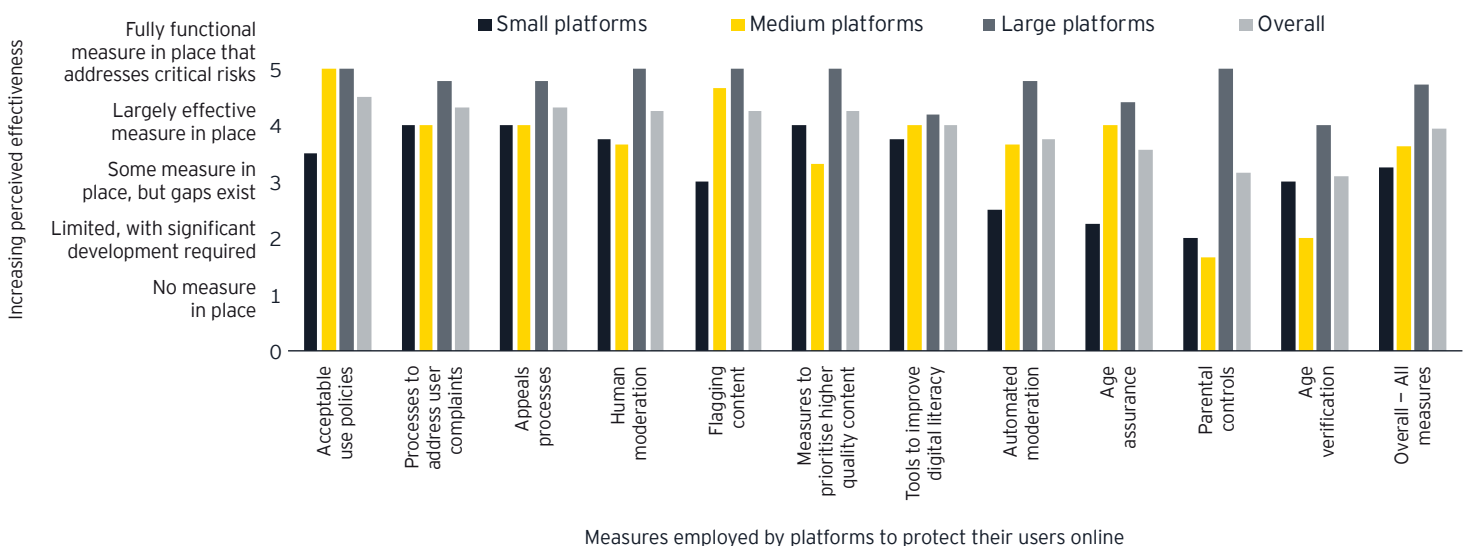
All of the platforms included in our research have measures in place to protect their users from harmful content, but the perceived effectiveness of these measures varied significantly. While the largest platforms had a suite of measures that they considered to be largely or fully effective in protecting users, some of the smallest platforms had more limited measures in place and were more likely to believe that the specific measures we asked about required significant development to make them effective.

Among the platforms included in our research, the most commonly implemented measure overall was acceptable use policies, which large and medium platforms generally

considered to be fully functional at addressing critical risks. However, the permissiveness of these policies varies significantly by platform, and it is against these platform-specific policies that all content review, removal and appeal decisions are made.

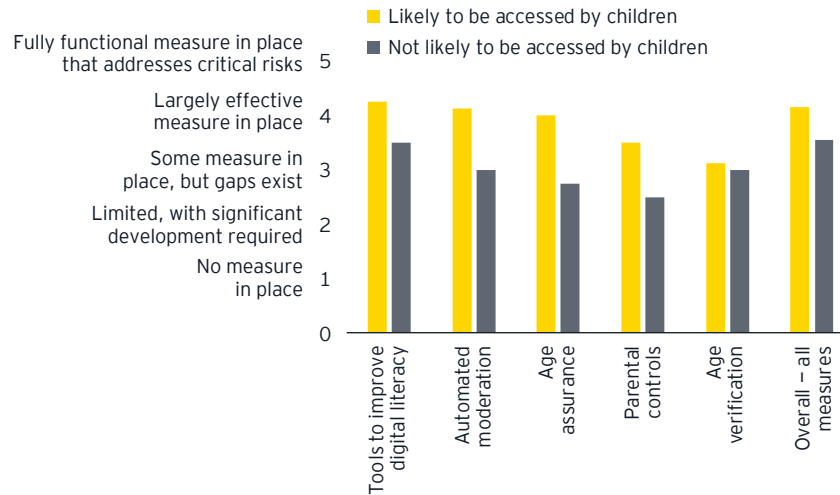
Technological and algorithm-driven solutions, such as automated moderation and age assurance measures, were perceived to be less effective among small platforms compared to large platforms. Large platforms generally have more resource to develop these solutions in house, which can be more cost effective for platforms of scale with the relevant engineering expertise.

Platforms' perceived effectiveness of measures to protect users online by platform size



Platforms that consider themselves likely to be accessed by children tended to report having more effective measures in place to protect their users from harmful content online. In particular, platforms that consider themselves likely to be accessed by children were more likely to have effective automated moderation, age assurance and parental controls in place compared to platforms that did not consider themselves likely to be accessed by children. There was less difference in the effectiveness of age verification measures in place.

Platforms' perceived effectiveness of selected measures by platform type




Overall, the platforms included in our research considered the measures that they had in place to be largely effective at protecting most users from harmful content online. We noted a potential inconsistency relating to the small platforms included in our research. Small platforms considered that the suite of measures they had in place were largely effective at protecting users overall. At the same time, small platforms reported that significant development would be required to make several of the individual measures we asked about effective, including flagging content, automated moderation, age assurance, parental controls and age verification.

Our research suggests that, rather than focusing on individual measures, risk assessments need to consider the suite of measures a platform has in place in light of the risks specific to the platform. Some measures will be more important to some platforms than others, for example depending on the type of content they host and whether they are likely to be accessed by children.

On the next three pages, we set out our assessment of the policy, prevention, detection and review measures that platforms have in place, based on the interviews and surveys we conducted with platforms.

Note: We refer to 'small' platforms as those having less than 100k unique global users, 'medium' platforms as those having between 100k and 10m unique global users, and 'large' platforms as having more than 10m unique global users.

Policy



Policy measures include terms and conditions, acceptable use policies and community guidelines. They are used across most platforms, and explained when a user first subscribes.

Terms and conditions, acceptable use policies and community guidelines

Most platforms consider user terms and conditions the foundation of their measures to protect users online.

Users generally have to accept terms and conditions when signing up to join a platform. These conditions include guidelines on what type of content is permissible under the platforms guidelines.

All platforms we spoke to explicitly stated that illegal content is banned. However, there is more uncertainty with potentially harmful but legal content, where judgement is required to determine whether content should be removed.

“

Community guidelines form the basis of how we self-regulate our platform.

To address this, we found some platforms allowed content creators or originators to define community guidelines, which define the rules of specific subsections within platforms (e.g. only certain conversation topics are allowed in a specific chatroom). Our research highlighted that these tended to be informal measures enforced by volunteers. We also found platforms each take different approaches to how permissive their terms, conditions or similar guidelines are. Some platforms appeared to have more permissive and community led guidelines, whereas others were more explicit in prohibiting certain types of content or behaviour.

“

Our user guidelines are designed to protect our users without inhibiting the right to free speech.


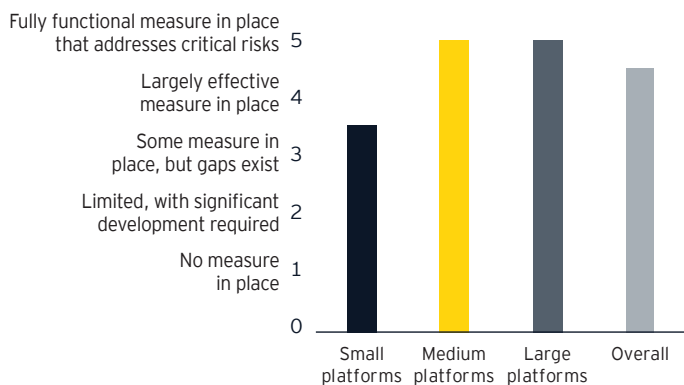
Platforms regularly update their user guidelines in response to what their users consider to be acceptable behaviour on the platform. We also found that user perception of acceptable behaviour changes, and platforms need to adapt constantly to effectively address risks.

“

Analysis of our user complaints helped us understand that misinformation has become an area of concern in recent years.

Our survey suggested that medium and large platforms had fully functional user terms and conditions or similar in place. However, small platforms suggested some gaps exist and improvement was required before policies would effectively address critical risks.

Platforms' perceived effectiveness of policy measures by platform size



Prevention

Prevention measures relate to stopping harmful content reaching users, either by preventing certain users from uploading content or by stopping harmful content from spreading. Measures are applied by platforms depending on the risks the platform perceives to be most relevant, given the type of content it hosts and the demographic of its user base.



Age verification

Age verification measures, where users are required to provide government ID, were sporadically applied across the platforms included in our research. Small and medium sized platforms were more likely to report having no or limited age verification measures in place compared to large platforms.

There is a trade-off between privacy and safety when it comes to verifying age online. This trade-off is more pronounced for adult sites, where users are likely to be more reluctant to provide ID. One adult platform discussed the potential for competitive distortion if the requirement for age verification was applied inconsistently, as users would be more likely to seek out non-compliant sites.

Some platforms discussed the challenges of verifying ages between 13 and 18, as these users are less likely to have formal ID. Age verification, along with age assurance, was mentioned by some platforms as an area where they may need to do more if they are in scope for upcoming UK regulation.

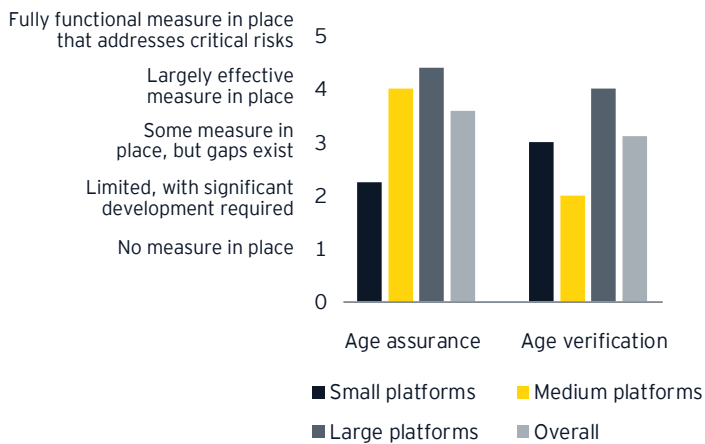
Age assurance

Age assurance measures vary significantly in sophistication. Some platforms have a checkbox where users self-certify that they are over 18. Others, typically large platforms or some platforms that are primarily used by children, have algorithmic-based checks that estimate the user's age based on their browsing behaviour or by scanning the user's face.

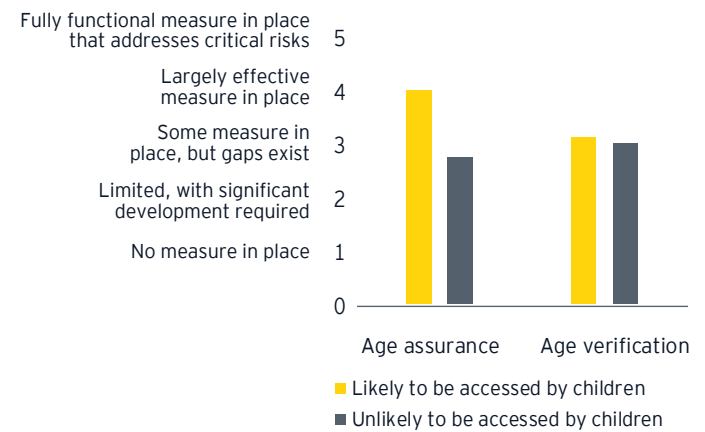
Generally, the perceived effectiveness of age assurance improves with platform size, with large platforms reporting they have largely effective measures in place. Platforms that are likely to be accessed by children were also more likely to report having effective age assurance mechanisms in place.

Among platforms with algorithmic age assurance measures, large platforms told us that they developed the measures in house, whereas smaller platforms tended to buy off-the-shelf software to avoid in-house development costs.

Platforms' perceived effectiveness of age measures by platform size



Platforms' perceived effectiveness of age measures by platform type



Pre-upload moderation and safety by design

The majority of platforms we spoke to do not undertake pre-upload content review to determine whether content complies with their acceptable use policies. Given the scale of content uploaded, this measure was seen to be impractical by some platforms. One adult site we spoke to said that it carried out pre-upload content review.

We also note that some platforms have created child friendly versions of their main sites, limited to content that has previously been reviewed by moderators and deemed suitable for children.

Generally, platforms seek to put the onus on the content uploader to ensure content is compliant, sometimes by requiring the uploader to check a box to confirm the content is compliant. There is a question around how thoroughly content uploaders engage with these terms and conditions.

Parental controls

A range of parental controls are available which, if enabled by the user, block access to websites that have added 'Restricted to Adults' tags. Pornography sites tend to be tagged as 'Restricted to Adults'. Parental controls are available for both desktop use and for mobile use, where they can block age inappropriate sites where a request for access has been made from a mobile phone registered to a child.

Parental controls have some limitations. While sites tagged as 'Restricted to Adults' can be blocked, explicit sexual or violent content can also be accessed on other social media sites where users are not required to be over 18, or where users self-certify their age.

Some platforms have in-app parental controls that allow parents to limit the functionality of the app for children, or that allow parents to track children's use of the platform.

Detection



All of the platforms included in our research have some measures to detect content that may conflict with the platform's acceptable use policies, but the sophistication of these measures varies. Due to the volume of content uploaded, many platforms focus their efforts on measures to detect potentially violative content after it is uploaded.

User flagging

Most platforms included in our research had measures in place to allow users to flag content for review if the user believes that content violates the platform's policies. Small platforms were more likely to think that their flagging measures required significant development to become effective.

As well as measures that allow general users to flag content for review, some platforms have 'trusted flagger' systems in place, where certain organisations (such as NGOs or child protection bodies) are designated as trusted flaggers, and content flagged by them is prioritised for review and removal.

Some platforms report a high number of 'false flags', where content flagged by users is ultimately not found to have contravened policy following review by moderators.

Automated moderation and machine learning

Automated moderation plays a key role in detecting child sexual abuse material (CSAM) and can automatically remove known CSAM and illegal content that is hashed on a database.

Platforms also use algorithms to detect a broader range of violative content. Algorithms can proactively flag videos for review more quickly than is possible when solely relying on user reports or human moderators. In its published transparency report, YouTube stated that 94% of videos removed in Q3 2020 were first detected by algorithm.

While some of the largest platforms have invested significantly in algorithmic moderation, there is a general sense that the technology is in its infancy for moderation of video content, requiring further development to be fully effective. As such, aside from CSAM and spam, most videos detected by algorithms must be reviewed by human moderators before removal.

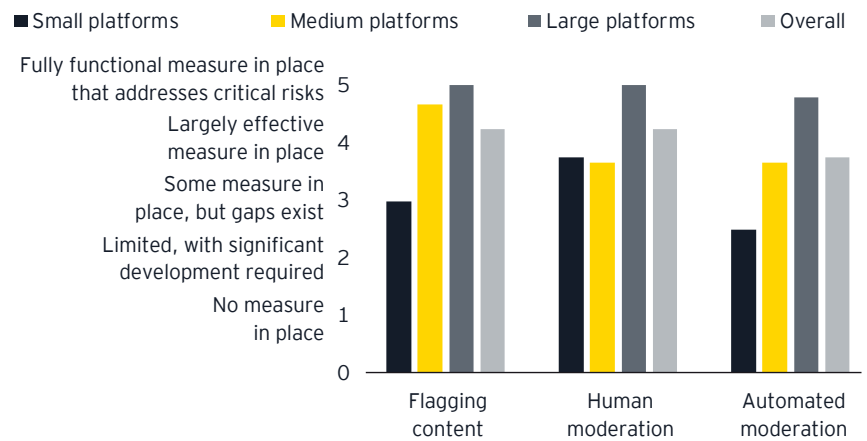
Human moderation

While automated moderation can play a role in detecting potentially violative content, most platforms rely heavily on human moderators to review videos to confirm whether they violate the platform's policy before the videos can be removed.

Human moderators can be outsourced or in-house employees of the platforms. Due to the volume of videos uploaded, effective human content moderation is resource intensive. 24/7 teams of moderators are required in all of the markets where the platform has a large user base to respond to content flagged by users or detected by algorithms in an appropriate time.

Human moderation necessarily relies on a degree of judgement, but can take account of the context of the video more effectively than algorithms are currently able to.

Platforms' perceived effectiveness of detection measures by platform size



Review

The majority of platforms included in our research were confident in the effectiveness of the process they had in place to address user complaints and to allow users to appeal decisions taken. However, smaller platforms we spoke to had concerns about the practicality of transparency reporting.

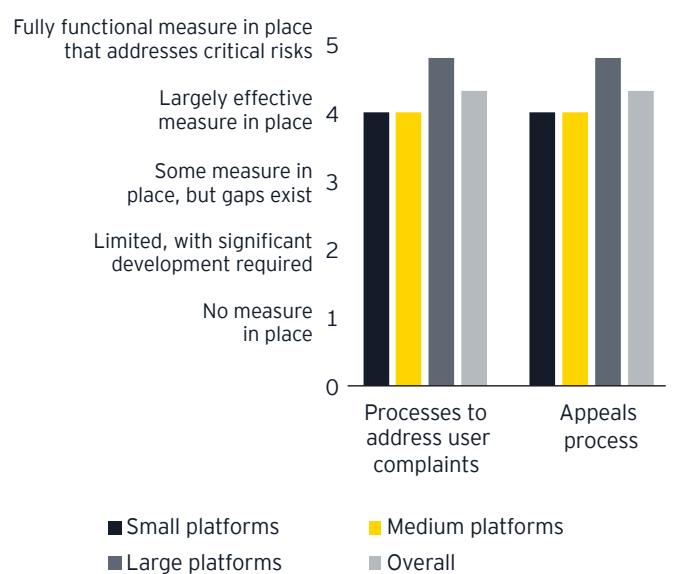
Process to address user complaints and appeals

Most platforms included in our research felt confident that they had processes in place to effectively address user complaints and allow users to appeal decisions taken. Some large platforms have expert moderation teams that specialise in different types of harmful content, and these teams play an active role in appeals.

Transparency reporting

Transparency reports published by providers disclose a range of information relating to content that has been reported to, and reviewed by, the platform's moderators. Transparency reporting requires significant investment and resource to develop the infrastructure to track and report data in a centralised way. Transparency reporting can be particularly challenging for some platforms who may not routinely collect this type of data. If all platforms are required to produce transparency or compliance reports, some believe they may need to redirect investment from operations to compliance reporting.

Platforms' perceived effectiveness of review measures



Our research suggests the types of cost that platforms incur are similar, but the scale and focus of expenditure varies



The main drivers of costs for platforms when implementing measures to protect users online are staff costs, technology development and infrastructure costs. Maintaining large teams of human moderators is a significant expense for some platforms, and brings associated costs such as investment in well-being packages and counselling support for moderators watching particularly upsetting content.

“

To deal with the volume of traffic on our platform, we need thousands of human content moderators in different parts of the world. Staff turnover is high because of the nature of the role, increasing our costs.

“

We invest in wellbeing packages and counselling to support moderators, as well as technology solutions to blur the most upsetting parts of videos being reviewed.

Costs generally scale with traffic: sites with more content require more moderators to effectively protect users from harm. Third-party technology solutions, for example those providing age assurance or age verification checks, are generally priced per check. However, there are substantial volume discounts available for large platforms, reducing the cost per user materially.

The exception to the above is in-house technology development, which is generally a fixed cost. For this

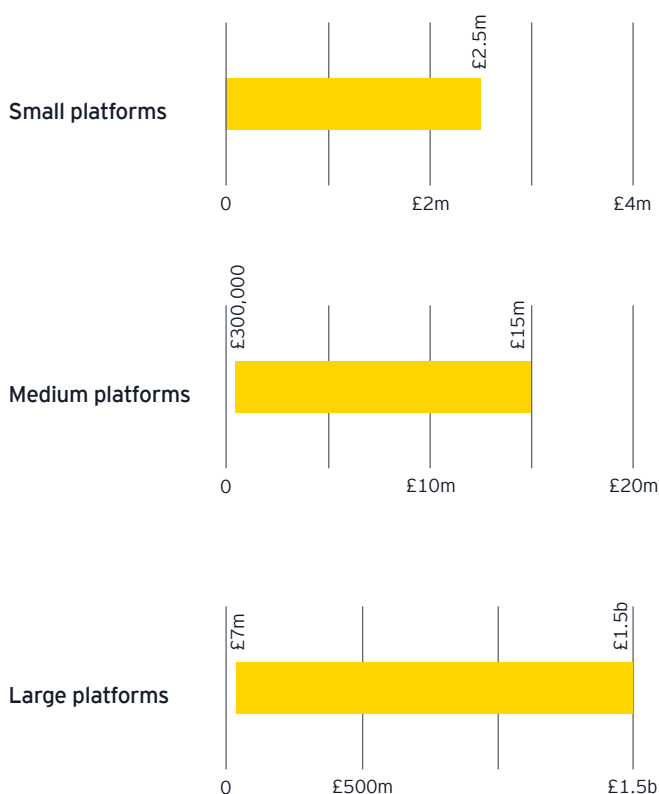
reason, large platforms with teams of engineers and in-house expertise will often seek to develop technological solutions themselves rather than purchasing from third parties. For small platforms with more limited resources, developing solutions in house may not be practical.

“

The costs of implementing off-the-shelf technology can be expensive for smaller platforms, as we don't benefit from economies of scale like larger players.

Total expenditure⁴

Total annual expenditure on measures to protect users online



Sources: Mark Zuckerberg Facebook post, February 2019.

⁴ Analysis of costs relates to total costs for measures to protect a platform's users, including users outside the UK.

Cost per user

Large platforms benefit from significant economies of scale when implementing measures to protect users online. Cost per user ranged from £0.25 to £0.50 for large platforms included in our research, with evidence suggesting the cost per user for some of the largest platforms may be materially lower.

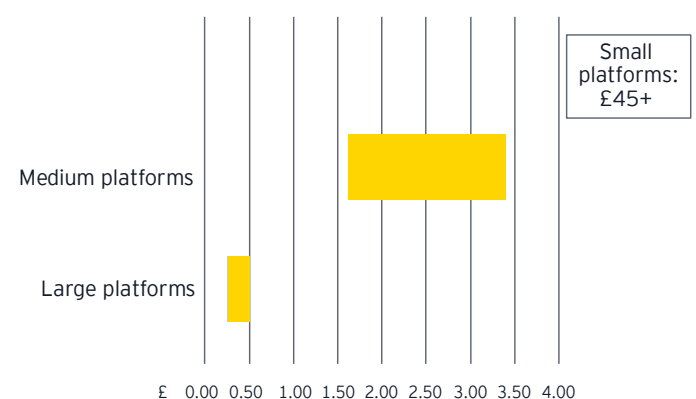
The cost per user for small platforms is likely to vary materially depending on the measures they have in place and on their lifecycle stage (e.g. loss-making start up vs. established small business). Small platforms included in our research spent over £45 per user on online safety measures.

Total annual expenditure on measures to protect users from harmful content varies significantly between different platforms, with large platforms spending substantially more than small platforms.

Total annual costs ranged from hundreds of pounds for the very smallest platforms to over £1.5bn for the largest platforms included in our research.

Some large platforms may spend significantly more than this. For example, Facebook publicly announced that it would spend more than \$3.7bn (c.£2.65bn) on safety and security on the platform in 2019.

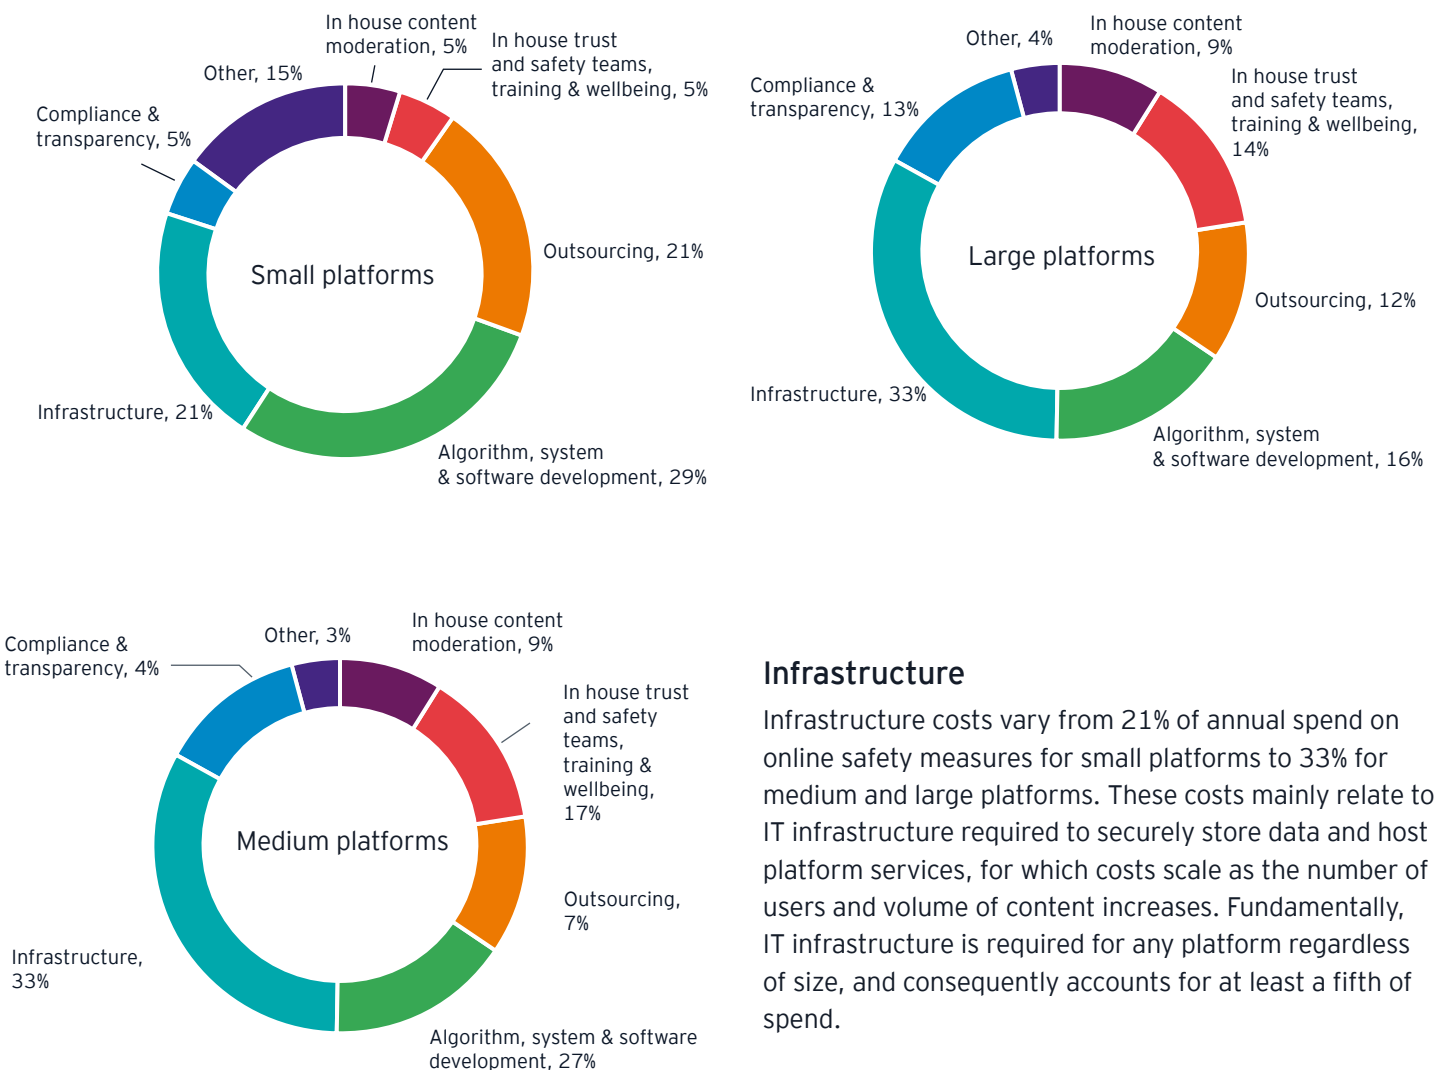
Cost per user by platform size



Investment allocation varies by platform size, with funding limitations at smaller platforms increasing the prevalence of outsourcing

Our research has demonstrated that different size platforms adopt different approaches to protecting their users from harmful content online. Because of this, there are some differences in the measures that platforms invest in and the types of costs they incur.

Contributions to annual spend on online safety measures by platform size



Infrastructure

Infrastructure costs vary from 21% of annual spend on online safety measures for small platforms to 33% for medium and large platforms. These costs mainly relate to IT infrastructure required to securely store data and host platform services, for which costs scale as the number of users and volume of content increases. Fundamentally, IT infrastructure is required for any platform regardless of size, and consequently accounts for at least a fifth of spend.

In-house content moderation

Medium and large platforms invest 9% of their annual spend in in-house content moderation. In comparison, small platforms only invest 5% of their annual spend in the same activity. Our research suggests recruiting and training in-house content moderators can take up to two years, and small platforms may not have the resources to prioritise this activity considering the lead times required to make it effective.

Algorithm, system and software development

Platforms invest 16% to 29% of their annual spend on these activities, including enhancing automated moderation capabilities. For small platforms, spend on this activity is up to £200k per year, in comparison to >£100m for large platforms. The much lower annual spend by small platforms also correlates to the sophistication of their solutions, with our survey suggesting significant development is still required to make them largely effective.

In-house trust and safety, training and wellbeing


At small platforms, these activities account for 5% of annual spend, in comparison to 17% and 14% for medium and large platforms. A majority of the costs for these activities relate to employing in-house trust and safety experts, where small platforms may seek to operate lean business models with skeleton teams. In contrast, medium and large platforms have significantly larger in-house teams, and subsequently training and wellbeing costs, which means these activities contribute much higher proportions of overall annual spend.

Outsourcing

Our analysis has found that small platforms are more likely to outsource measures (including technological and legal support), contributing 21% of annual spend. In contrast, medium and large platforms only invest 7% to 12% of their annual spend in outsourcing, with their priority being to develop in-house solutions where possible.

Other costs

For small platforms, other costs contribute 15% of total annual spend, and relate mainly to end-user support.



Regulation and legislation is one of several factors that drive investment in protecting users online



Our research suggests several factors drive investment in protecting users from harmful content online. Those drivers include:

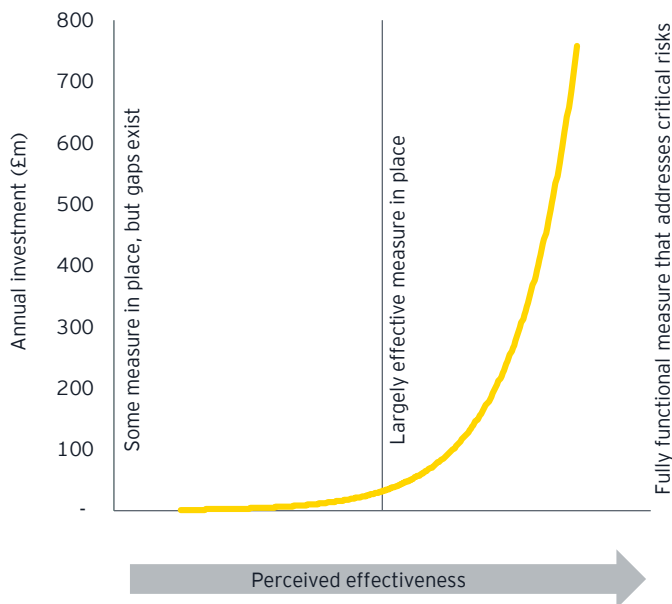
Maintaining brand integrity: Platforms seek to ensure the integrity of their brand is maintained by continuing to improve user safety. High-profile breaches of user safety could have detrimental impacts on brand perception, and subsequently user numbers.

Evolving risks: Our research has identified that risks to online safety are not consistent, and evolve over time. 'Bad actors' find different ways of causing online harms and social standards change, meaning that platforms need to continuously invest to protect their users.

Potential legislation and regulation: Although legislation and regulation in relation to the obligations of online platforms is still evolving, many platforms are taking a pre-emptive approach to ensuring processes and systems are in place to comply with future requirements.

Our research highlighted that no specific factor is considered more of a priority than others. Instead, future investment is prioritised based on a combination of the risks the platform faces, resource constraints, and the platform's longer-term plans.

Annual investment vs. Platforms' perceived effectiveness of measures



The investment required to make measures fully functional

Our indicative analysis, based on our survey of 12 platforms of varying size, found significant investment is required to improve the effectiveness of measures employed to protect users from harmful content.

When analysing the perceived effectiveness of measures against the amount invested, we found investment increases exponentially as the effectiveness of measures improves, and there was a diminishing benefit to incremental investment.

Because of this relationship, most platforms suggested they take a risk-based approach to assessing which specific types of harm should be prioritised for action on their platform.

The impact of the AVMSD on investment in protecting users from harmful content online

Most platforms included in our research suggested the AVMSD would not lead to any incremental investment requirements. Other factors, as described above, drove a majority of their future investment decisions.

One platform did suggest the AVMSD would lead to the reprioritisation of existing investment plans. It suggested it would bring forward planned age verification and age assurance investment to improve its capabilities and

ensure compliance, but these improvements would be funded by delaying other planned investment.

Only one platform suggested it may invest in age verification and age assurance technologies as a result of the AVMSD, which it was unlikely to have done otherwise. From the research undertaken, this represented the only instance of incremental investment caused by the implementation of the AVMSD.

Platforms of all sizes think that there are a range of barriers to effectively protecting users online

Overall, complexity or understanding of regulatory requirements was perceived to be the most significant barrier to effectively protecting users online by the platforms included in our research. A particular concern raised by the platforms we spoke to related to the need for flexibility.

“ We have limited resources and we need the flexibility to decide how we will respond to content that is flagged by users so we are able to prioritise the most harmful content.

“ Not all measures are relevant to every platform. We know the risks that are most relevant to our own platform.

Small platforms were more likely to consider technological limitations and cost to be barriers to protecting users online. Small platforms often do not have the engineering resource that large platforms have to develop technology solutions in-house, and may also not have the resource to have large teams of human content moderators. Cost was perceived to be a less significant barrier by larger platforms included in our research compared to smaller platforms.

“ Technology is the biggest barrier. To be effective, many of these measures need to be automated, but the technology is years away from where it needs to be.

Lack of collaboration was generally perceived to be a less significant barrier. However, some platforms we spoke to noted the challenge of accessing different databases of CSAM material held by different organisations in different jurisdictions. Others felt that, while collaboration between platforms on CSAM material and terrorist content was well-developed, there could be more collaboration between platforms on other types of harm. For example, these platforms felt there could be more sharing of other types of harmful content that has been removed in order to assist with detection of that content on other platforms. One platform noted that collaboration is much more common in other industries than it is in the social media and online video sector.

“ There are several different lists of CSAM held by different organisations, with little collaboration. It can be prohibitively expensive for small platforms based outside the US to access some of these lists.

Loss of competitiveness was perceived to be less of a barrier by medium and large platforms, which are more likely to have the resources to invest in developing innovative technologies in-house. For smaller platforms, loss of competitiveness was perceived to be more of a

concern. One platform noted that the very largest platforms have the resources to develop the best algorithms, and that there is a question about how these technologies can be made accessible for smaller platforms. Another platform noted concerns around loss of competitiveness if regulation is applied inconsistently across the industry, particularly in relation to age verification.

Finally, platforms we spoke to raised the challenge of adaptive ‘bad actors’, who adapt to undermine new safety measures implemented. These adaptive bad actors, coupled with the long lead time for technology development, can undermine efforts to protect users.

“ It takes, on average, two years to develop, test and integrate technological solutions on our platform. In that time, bad actors have found new ways of gaming the system.

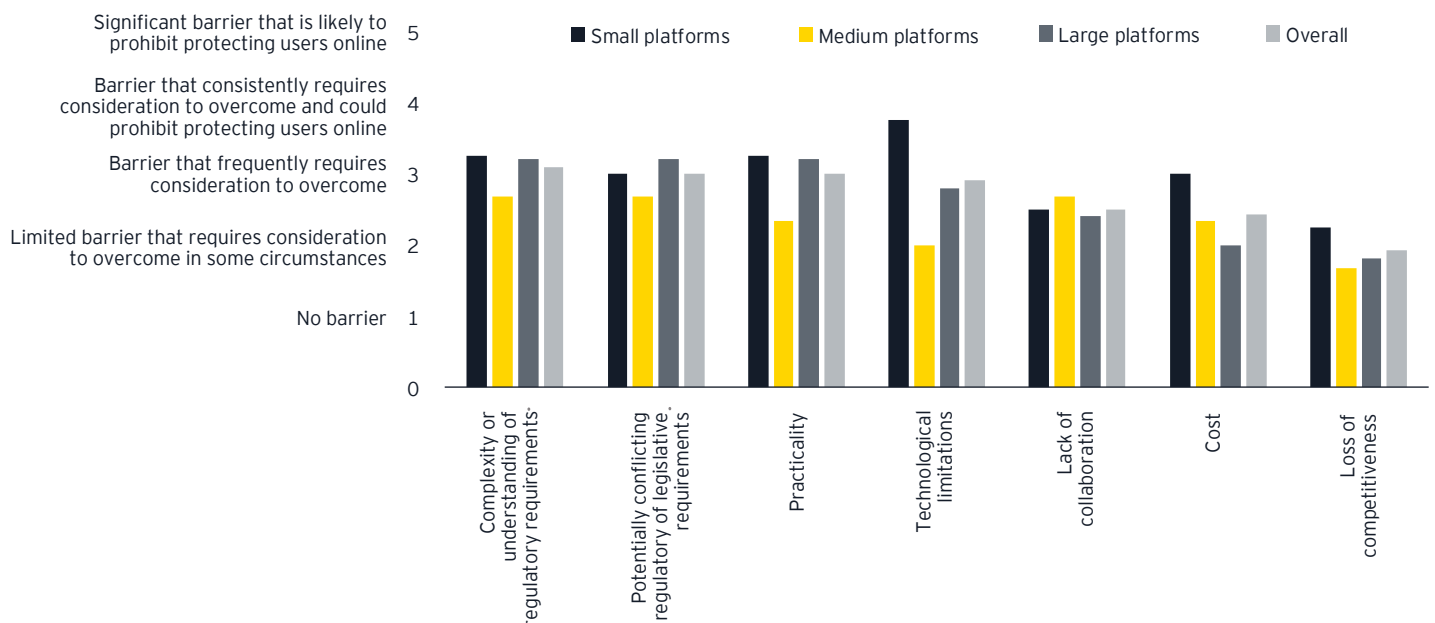
The role of media literacy

Platforms told us that changing user expectations are a critical factor influencing the steps platforms take to protect their users online. Because of this, as users’ media literacy improves, their expectations of the measures platforms should have in place are likely to increase, incentivising the platforms to further invest in online safety. The measures platforms implement play a crucial role in protecting users online. Improving media literacy will also help create a safe environment online.

Improving media literacy will help users act appropriately, consider the safety of themselves and other users of the platforms, and critically evaluate information they encounter online.

As user literacy improves, their expectations of the measures platforms should have in place is likely to increase, incentivising the platforms to further invest in online safety.

Perceived barriers to protecting users online by platform size



Findings

Measures and costs

1

Different platforms take different approaches to protecting users online, driven by factors including the size of the platform, the resources it has to invest in online safety, its assessment of the risks specific to the platform, the perception of its brand, and, in some cases, competitive considerations.

2

Generally, the largest platforms have the most confidence in the measures they have in place to protect users online. With some exceptions, small platforms are more likely to state that the specific measures we asked about require significant development to effectively protect users online.

3

Content moderation can be resource intensive, with technological solutions, such as algorithmic moderation, still in their early stages of development. Effective, responsive content moderation requires large teams of human moderators.

4

The costs of protecting users online varies significantly by scale. Large platforms tend to have the resource to develop technology solutions in-house, which are largely fixed costs. The cost of third-party solutions scale by traffic, but with significant volume discounts making these solutions more cost effective on a per user basis for large platforms compared to small platforms.

5

There are diminishing improvements from incremental investment in online safety measures, and platforms take a risk-based approach to prioritise investment to allocate limited resources.

6

Small platforms are much more likely than large platforms to view technological limitations and cost as barriers to protecting users. For the large platforms included in our research, cost was considered to be only a limited barrier to protecting users.

5

We have assessed the user experience of platforms with video-sharing capabilities



Our primary consumer research has focused on understanding the user experience of platforms with video-sharing capabilities. We have focused on users' experiences of harm, expectations for safety online, and willingness to engage with measures to protect them. In this context, we have administered two surveys covering three sample groups: adults aged 18+, children aged 6-12, and children 13-17. We have assessed the user experience of these platforms across each of the following topics:

01

How do people use platforms with video-sharing capabilities?

03

What do people think about harmful content online?

02

What is the user experience of harmful content online?

04

How do people respond to measures taken by platforms to keep them safe online?

The sample sizes for the groups included in our survey are as follows: adults aged 18+ (1,000 respondents), children aged 13-17 (500 respondents), children aged 6-12 (500 respondents). All respondents are UK-based users.

01 | How do people use platforms with video-sharing capabilities?

Different age groups find the content they want to watch in different ways

44% of adults said the main way they find videos is using search functions, either using a search engine or searching directly within the platform itself. Recommendation algorithms also play an important role: when asked to list all the ways that they find videos online, one in four adults say they browse videos recommended by the platform itself.

In contrast, just 27% of 13-17 year olds said the main way they find videos is using search functions, either using a search engine or searching directly within the platform itself. When asked to list all the ways that they find videos online, one in three 13-17 years olds say they browse videos recommended by the platform itself.

More popular methods among 13-17 year olds were viewing content shared by friends or social media contacts (50%), and scrolling through videos that play automatically, such as on TikTok (37%). Younger age groups are also more likely to search for specific content providers or search for content on channels they subscribe to compared to older groups.

The measures we have reviewed as part of this report have limited focus on recommendation algorithms, despite this being a common way in which video content is accessed. Platforms and policymakers should consider how video content is accessed when assessing the risks that users face.



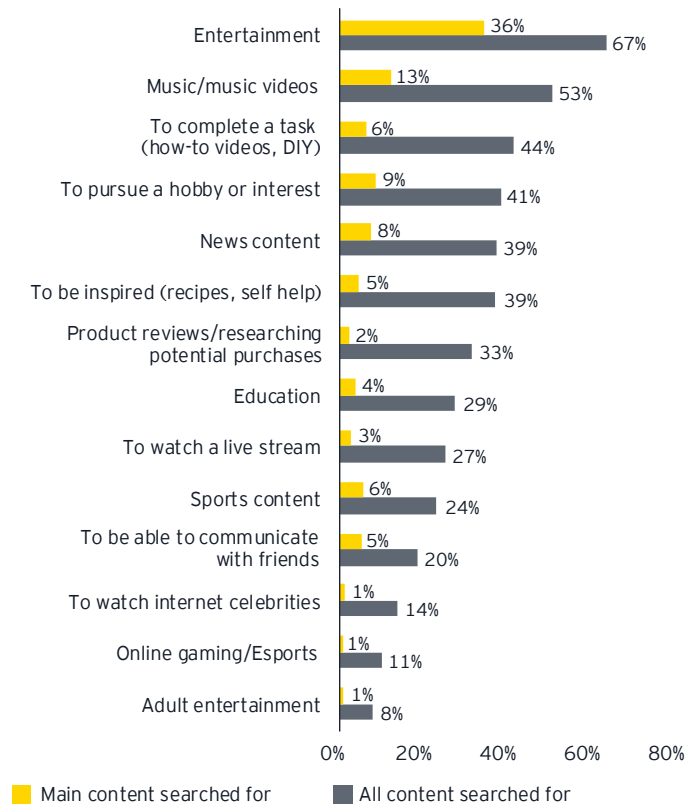
Entertainment and music are the most popular types of videos online

While most searches are for music or entertainment, content related to how to complete tasks, hobbies and interests, and self help are also highly popular. Additionally, 40% of users search for news content on platforms with video-sharing capabilities. Exposure to misinformation on news-style content could be a risk for users.

For 13-17 year olds, content related to online gaming and e-sports is popular, with 35% of users searching for this content and 11% stating it is the main content they search for.

Computer games are popular among those aged 6-12, with 28% stating it was the content they enjoyed the most, with 'funny videos' ranking second with 17% of users stating this option.

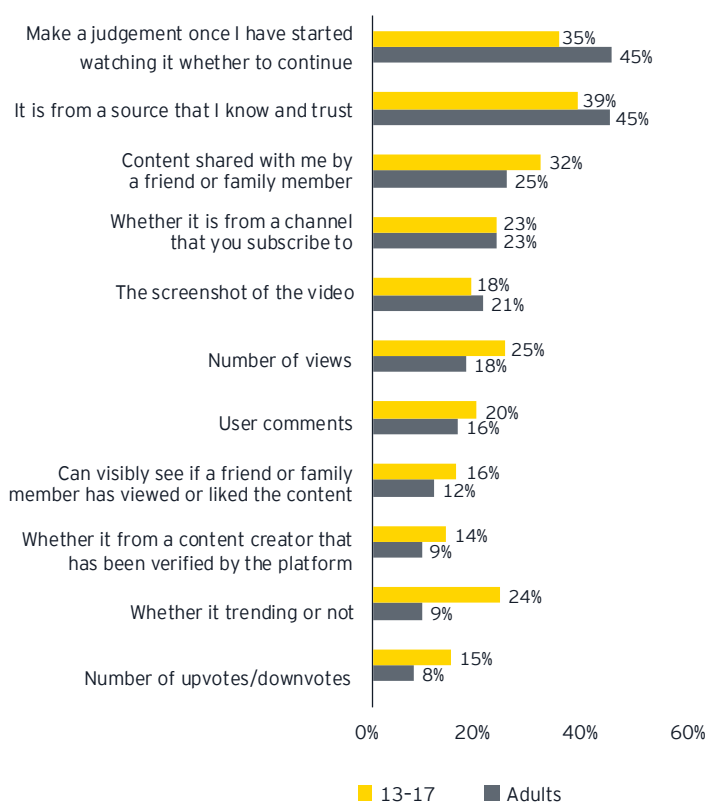
Main/all content searched for in the UK, adults 18+



Trust in the source of content is important

Adults are just as likely to make a decision on whether to continue watching a video based on how much they trust the source of content as they are to make this judgement once they have started watching the content. 45% of adults and 35% of 13-17 year olds make a judgement on the appropriateness of the video after they have started watching it. The potential for accidental exposure to harmful content online should therefore be considered by platforms and policymakers.

Important factors when deciding whether to watch a video



Most young children watch videos without supervision and are attracted to sites designed for older children

Only 12% of young users watch with an adult family member present, with 66% watching videos by themselves. 54% agree with the statement: "I am allowed to use video apps or websites that are meant for older children, as long as I am careful", while 50% agree with the statement: "Video apps or websites that are meant for older children are more fun than those which are meant for people my age". The way that children use platforms with video-sharing capabilities suggests that children may be particularly at risk of being exposed to harmful content online.

49% of those aged **6-12** are constrained to sites chosen by their parents or guardians

People upload content for a range of reasons

54% of uploaders do so to share content with friends and family, and 40% to share content associated with a hobby or interest.

14% share content to support an online business they run, highlighting the increasing role of platforms in supporting businesses, and 15% to share content for a charity or not-for-profit organisation.

Younger users are more likely to upload content: 48% of those aged 13-17 upload and share content, with 24% of these sharing content in the hope of 'going viral'.



Key considerations for policymaking: One of the most common approaches to deciding whether to watch a video is to make a judgement once the video has started playing. This approach could risk users being accidentally exposed to harmful content. Additionally, most young children watch videos without supervision and are attracted to sites designed for older children, highlighting the need for platforms to consider the potential for their content to be accessed by children.

02 | What is the user experience of harmful content online?

Stated exposure to harmful content is high

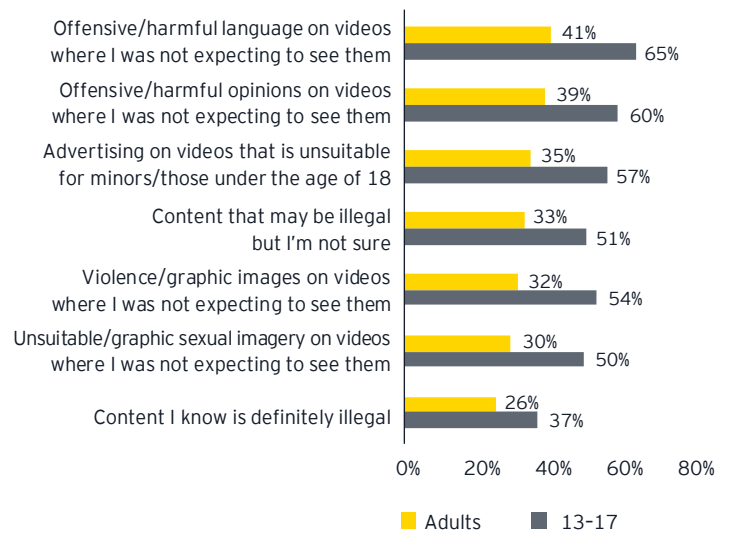
Stated exposure to harmful content is much higher among 13-17 year olds compared to adults across all of the categories of content we asked about.

Accidental exposure to offensive or harmful language on videos was the most commonly experienced type of harm, with 41% of adults and 65% of 13-17 year olds saying they have experienced this type of content in the last three months.

A third of adults and half of 13-17 year olds have seen content online in the past three months that they think may have been illegal, while 26% of adults and 37% of 13-17 year olds say they have been exposed to content that they believe is illegal in the past three months.

The platforms included in our research reported having largely effective measures in place to protect users from harmful content online. However, our consumer research showed the scale of harmful content accessed online, highlighting the risks for users. Our analysis suggests that significant incremental investment would be required to materially improve the effectiveness of measures employed by platforms.

% agreeing they have seen this type of content in the past three months

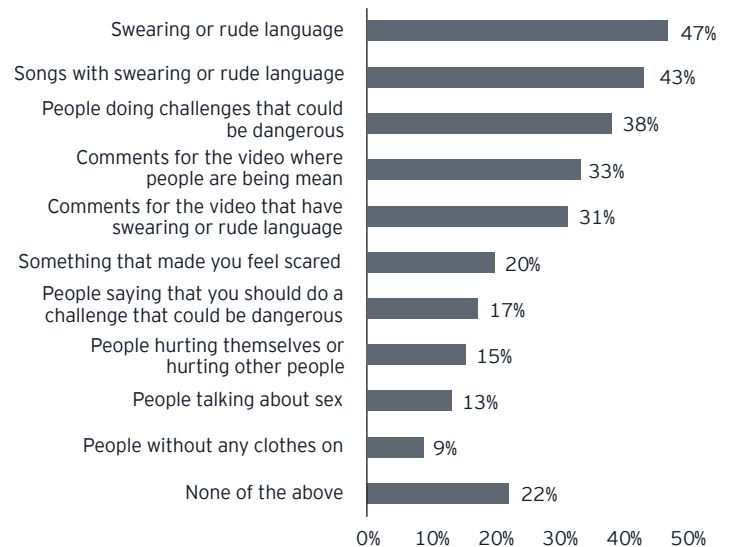


78% of 6-12 year olds have experienced some type of harmful content online

There is good awareness of online harm in this age group, with 94% agreeing **“there can be bad things on websites that can be accidentally seen if not careful”**. 89% of children aged 6-12 say that they would tell a parent or guardian if they saw a video that upset them. The biggest area of harm exposure was around bad language, with many experiencing this in music.

Online pranks and challenges are popular content with this age group, with 7% and 6% respectively saying it was their favourite content. However, 38% of respondents highlighted seeing challenges that could be dangerous and 17% said that they felt encouraged to do a challenge that could be dangerous.

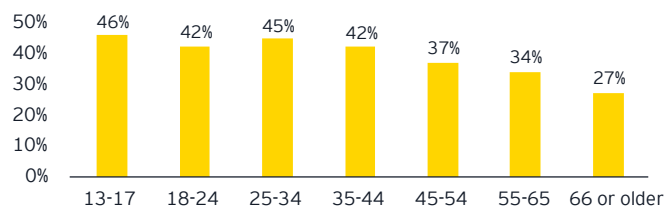
Harm exposure in 6-12 year olds (% ever experiencing)




Younger age groups are more concerned about accidental exposure to harmful content

Many people are concerned about accidental exposure to harmful content. This concern is less pronounced in older age groups, with only 27% of those aged 66 or older reporting that they are concerned about accidental exposure to harmful content online.

% concerned about accidental exposure to harmful content



Key considerations for policymaking: Exposure to harmful content on platforms with video-sharing capabilities is widespread across all age groups but particularly pronounced among under 18s. Some types of content could be more immediately harmful than others. For example, some 6-12 years old have been exposed to nudity, videos of people hurting themselves or others, and people encouraging them to undertake a dangerous challenge, while half of 13-17 years have unexpectedly seen unsuitable or graphic sexual imagery in the past three months.



03 | What do people think about harmful content online?

Some people struggle to identify truthful content

For most users, trust is important. 73% of those aged 13-17 and 69% of adults agree that they “would only be happy to watch video content on a video sharing platform that I know and trust”. This sentiment is more pronounced in older age groups: 75% of those over 55 agree.

However, some people struggle to identify what is and isn't true online. 30% of adults agreed with the statement: “I find it difficult to identify what content is true and what is false online”. This trend is much more pronounced in younger age groups with 41% of 13-17 year olds agreeing, compared to 19% of those aged 55 or over. It is unclear whether this trend reflects differing levels of awareness of misinformation online across different age groups. This could lead to people putting misplaced trust in harmful or inaccurate sites or sources.

There are differing views on what type of speech and behaviour is acceptable online

66% of adults agree that freedom of expression online includes “a duty to behave responsibly and respect other peoples' rights”. However, 22% think it means that “people should be free to express any opinion and say/show what content they want online without any interference or censorship from government bodies and/or the video sharing platforms”. This sentiment increased to 31% of the 13-17 age group.

This confusion may stem from people taking cues from the US, where freedom of speech laws are more permissive. 30% adults agreed that “the United States and the United Kingdom have similar laws and rules governing what people can say and do online”, while 30% disagreed and 40% did not have a view.

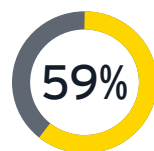
Differing views on what is acceptable online could contribute to differing attitudes to intervention by platforms

Understanding of freedom of speech can influence the way people respond to intervention by platforms. 19% of adults agree that “It's not right that video-sharing platforms are increasingly policing/censoring content that is shared on their sites”.

About 31% of people aged over 13 agreed that the “removal of high profile people from social media platforms for expressing their views is wrong/inconsistent with the principle of freedom of speech, even if it breaks the terms of service of the platform”.

A significant proportion are happy for harmful content to be ignored. 52% of 13-17 year olds and 43% of adults agreed with the statement: “If I accidentally see harmful or offensive (but legal) content online, I am happy to forget about it and move on”. However, 59% of UK adults with the statement: “Video-sharing platforms are not doing enough to keep illegal content off their platforms”.

Our research highlights confusion among some users as to the role of platforms, the rights of users and the extent to which platforms and government should intervene. It also highlights the need for platforms to adopt flexible measures to respond to the changing expectations of their users.



of UK adults agree that platforms are not doing enough to keep illegal content off their sites

Key considerations for policymaking: Most people tend to only watch videos on platforms they know and trust, but users struggle to understand what is and isn't true online, which could result in them putting misplaced trust in sources of misinformation. People's opinions on the kind of speech and behaviours acceptable online vary, which influences the way they respond to measures put in place by platforms to protect users online. The measures employed by platforms need to balance the expectations and priorities of different users.

04 | How do people respond to measures taken by platforms to keep them safe online?

Engagement with site information is low

The user experience of the effectiveness of platforms' acceptable use policies and terms and conditions is low.

Only 26% of adults state that they read the terms and conditions or terms of use on platforms with video-sharing capabilities to better understand the risk of seeing inappropriate or harmful content before using the site. 57% also stated that when they have read platforms' terms and conditions, they have found them complex and difficult to understand.

There appears to be a gap between users and platforms in terms of the perceived effectiveness of acceptable use policies. Overall, platforms included in our research perceived their acceptable use policies to be largely effective in protecting users online.

Opinions were mixed on the usefulness of different measures to help people evaluate the appropriateness of videos before viewing them

When asked which measure would most help users evaluate whether videos are safe and appropriate to watch, no clearly preferred measures emerged.

22% said that a detailed description of the content of the video (e.g. existence/severity of use of bad language, violence, nudity etc.) would help them better evaluate videos before viewing. Age appropriate ratings, based on ratings from those who have watched the videos (14%) or on a rating provided by the content uploader (12%), were also key interventions that would help adults evaluate

whether videos are safe and appropriate to watch. 12% stated they would be helped by a short video introduction summarising its contents and any disclaimers.

However, 11% of adults and 9% of children aged 13-17 said that none of the measures included in our survey (relating to improving information, providing descriptions of videos, including details of uploaders, and providing age appropriateness ratings) would help them better evaluate videos.

A large proportion of respondents were resistant to safety measures that introduce friction in the user experience

42% of UK adults agree with the statement:
"I prefer video platforms where there is no barrier between me and the video if I click on it to watch"

44% of UK adults agree with the statement:
"I would prefer to use a platform where I can remain anonymous and don't have to verify my ID"

39% of UK adults agree with the statement:
"I only look at the videos on video-sharing platforms, so wouldn't pay attention to any additional links or information on the site designed to stop exposure to online harm"

Key considerations for policymaking: Given the current lack of engagement with platforms' acceptable use policies, it is unclear whether increasing text-based information would improve online safety. The large proportion of respondents who say they prefer video platforms where there is no barrier between them and the video suggests that interventions that are integrated into the user experience in a way that minimises friction would be most effective in protecting users online.

Findings

Consumer research

1

The way users engage with platforms with video sharing capabilities means there is a risk of accidental exposure to harmful content online. Platforms need to consider the potential for their content to be accessed by children regardless of whether their content is intended for children or not.

2

The platforms included in our research reported having largely effective measures in place to protect users from harmful content online. However, our consumer research showed the scale of harmful content accessed online among all age groups, with exposure to harmful content on platforms with video-sharing capabilities particularly pronounced among under 18s. Our analysis suggests that significant incremental investment would be required to materially improve the effectiveness of measures employed by platforms.

3

People have differing attitudes to harmful content online and the appropriate response to such content. Our research highlights confusion among some users as to the role of platforms, the rights of users and the extent to which platforms and government should intervene. The measures employed by platforms need to balance the expectations and priorities of different users.

4

Many people do not engage with information on platforms' websites, and those that do engage struggle to understand the information available. A large proportion of respondents to our survey were resistant to safety measures that introduce friction in the user experience, suggesting that interventions that are integrated in a way that minimises friction would be most effective in protecting users online.



EY | Building a better working world

EY exists to build a better working world, helping to create long-term value for clients, people and society and build trust in the capital markets.

Enabled by data and technology, diverse EY teams in over 150 countries provide trust through assurance and help clients grow, transform and operate.

Working across assurance, consulting, law, strategy, tax and transactions, EY teams ask better questions to find new answers for the complex issues facing our world today.

EY refers to the global organization, and may refer to one or more, of the member firms of Ernst & Young Global Limited, each of which is a separate legal entity. Ernst & Young Global Limited, a UK company limited by guarantee, does not provide services to clients. Information about how EY collects and uses personal data and a description of the rights individuals have under data protection legislation are available via ey.com/privacy. EY member firms do not practice law where prohibited by local laws. For more information about our organization, please visit ey.com.

Ernst & Young LLP

The UK firm Ernst & Young LLP is a limited liability partnership registered in England and Wales with registered number OC300001 and is a member firm of Ernst & Young Global Limited.

Ernst & Young LLP, 1 More London Place, London, SE1 2AF.

© 2021 Ernst & Young LLP. Published in the UK.
All Rights Reserved.

EYUK-000140696.indd (UK) 07/21. Artwork by Creative London.

ED None



In line with EY's commitment to minimise its impact on the environment, this document has been printed on paper with a high recycled content.

Information in this publication is intended to provide only a general outline of the subjects covered. It should neither be regarded as comprehensive nor sufficient for making decisions, nor should it be used in place of professional advice. Ernst & Young LLP accepts no responsibility for any loss arising from any action taken or not taken by anyone using this material.

ey.com/uk